

PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁷ : G06T 15/40		A1	(11) International Publication Number: WO 00/28480
			(43) International Publication Date: 18 May 2000 (18.05.00)
(21) International Application Number: PCT/GB99/03704		(81) Designated States: JP, US, European patent (A1, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).	
(22) International Filing Date: 8 November 1999 (08.11.99)			
(30) Priority Data: 9824412.2 6 November 1998 (06.11.98) GB		Published With international search report.	
(71) Applicant (for all designated States except US): IMAGINATION TECHNOLOGIES LIMITED (GB/GB); Home Park Estate, Kings Langley, Hertfordshire WD4 8LZ (GB).			
(72) Inventor; and (75) Inventor/Applicant (for US only): FENNEY, Simon (GB/GB); 99 Watford Road, St. Albans AL2 3JY (GB).			
(74) Agent: ROBSON, Aidun, John; Riddle & Grose, 16 Theobalds Road, London WC1X 8PL (GB).			
(54) Title: SHADING 3-DIMENSIONAL COMPUTER GENERATED IMAGES			
(57) Abstract			
<p>An apparatus for shading 3-dimensional computer generated images representing each object in the image by a set of flat polygons. Vertex data (x, y) is supplied for each vertex of the polygon along with data which defines the orientation of the surface of the polygon. Edge data is computed in edge processors (6). The edge data is equivalent to data which defines a surface which is perpendicular to a viewpoint and which lies normally towards the polygon. For each pixel which is used to view the portion of the image in which the polygon is located, a depth value is derived for each of these surfaces and, when the depth values indicate that the polygon is visible at the pixel, it is shaded accordingly.</p>			

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SW	Switzerland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KR	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KZ	Republic of Korea	PT	Portugal		
CN	China	LC	Kazakhstan	RO	Romania		
CU	Cuba	LI	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LK	Liechtenstein	SD	Sudan		
DE	Germany	LR	Sri Lanka	SE	Sweden		
DK	Denmark		Liberia	SG	Singapore		
EE	Estonia						

WO 00/28480

PCT/GB99/03704

- 1 -

SHADING 3-DIMENSIONAL COMPUTER GENERATED IMAGES

This invention relates to the shading of 3-dimensional computer generated images and to a method and apparatus for performing this.

5 In our British Patent No. 2281682, there is described a 3-D rendering system for polygons in which each object in a scene to be viewed is defined as a set of surfaces which are infinite. Each elementary area of the screen in which an image is to be displayed has a ray projected
10 through it from a viewpoint into the 3-dimensional scene. The location of the intersection of the projected ray with each surface is then determined. From these intersections it is then possible to determine whether any intersected surface is visible at that elementary area. The
15 elementary area is then shaded for display in dependence on the result of the determination.

The system can be implemented in a pipeline type processor comprising a number of cells, each of which can perform an intersection calculation with a surface. Thus
20 a large number of surface intersections can be computed simultaneously. Each cell is loaded with a set of coefficients defining a surface for which it is to perform the intersection test.

A further improvement which is described in our UK
25 Patent Application No. 2298111 sub-divides the image plane into sub-regions or tiles. This proposes using a variable tile size and projecting a bounding box around complex objects. This is done by firstly determining the distribution of objects around the visible screen for
30 suitable tile sizes to be selected. The surfaces defining the various objects are then stored into one contiguous list. This avoids the need to store identical surfaces for each tile, as one object being made of many surfaces could be in a number of tiles. The tiles can then be
35 rendered in turn using the ray casting technique described

WO 00/28480

PCT/GB99/03704

- 2 -

above, one at a time rendering all objects within that tile. This is an efficient method because no effort needs to be made to render objects which are known not to be visible in a particular tile.

We have appreciated that the amount of processing can be reduced further if only data pertaining to portions of surfaces which are in fact visible is processed. Thus, in accordance with a preferred embodiment of the invention we provide a method for defining the edges of visible surfaces with planes which are perpendicular to the viewing direction.

In accordance with a second aspect of the invention, we have appreciated that rather than use a variable tile size, the processing may be optimised by using a regular tile size across the whole of the image plane wherein the tile boundaries may intercept with objects but with no edge clipping being necessary. A set of tiles can then be selected which define a bounding box for a particular object and, in order to render that particular object, only the tiles within that particular bounding box needs to be processed. A display list of the surfaces which fall within that tile is used to define objects within the bounding box.

A further improvement on this method discards the tiles within a bounding box which do not actually contain the object to be rendered.

Preferred embodiments of the invention will now be described in detail by way of example with reference to the accompanying drawings in which:

Figure 1 shows a graphical representation of a triangular surface for use in representing a portion of an object;

Figure 2 shows how the positive and negative sides of the surfaces are used to determine the visible portion of the triangle of Figure 1;

SUBSTITUTE SHEET (RULE 26)

WO 00/28480

PCT/GB99/03704

- 3 -

Figure 3 shows schematically the edge processors and surface processors used to shade the triangle of Figure 2;

Figure 4 shows schematically a screen with an object on it with the screen divided into an array of 25 tiles and with an object broken up into triangles in a conventional method;

Figure 5 shows schematically the object lists which are used in accordance with an embodiment of the invention;

Figure 6 shows a triangle within a bounding box of 12 tiles;

Figures 7a, b, c and d show a number of different triangles with different bounding boxes and the different numbers of tiles required to display them;

Figure 8 shows an optimised selection of tiles for the triangle in Figure 7d;

Figure 9 shows a table which illustrates the tests used to determine the tiles in Figure 8 which are not required to display the triangle; and

Figure 10 shows a rectangular set of tiles with a test point; and

Figure 11 shows a block diagram of the circuits used to generate bounding boxes.

In our British Patent No. 2281682, the rendering system summarised in the introduction of this specification is described. We have appreciated that any object can be modelled as a set of triangles. Thus, these would be the infinite surfaces which would be processed in that patent. In that patent, the edges of the objects would comprise the intersections of the infinite surfaces and the relative depths of forward and backward facing surfaces used to determine whether or not a particular surface was visible (if a backwards facing surface is closer than a forwards facing surface then neither is visible at a particular pixel).

SUBSTITUTE SHEET (RULE 26)

WO 00/28480

PCT/GB99/03704

- 4 -

We have appreciated that processing can be improved by defining the edges of triangles by infinite surfaces which are perpendicular to the viewing point. Thus, for a triangle, four surfaces are required, one for the face and three for the edges, one per edge.

Before a triangle can be rendered, it is necessary to calculate the equations for each surface. These are calculated in a polygon setup unit from vertex data supplied by the application software. The equation for a perpendicular edge surface between two vertices v_1 and v_2 , shown in Figure 1 and which are located at (x_1, y_1, z_1) and (x_2, y_2, z_2) , is defined by:

$$(y_2 - y_1)x + (x_1 - x_2)y + (x_2y_1 - y_2x_1)z = 0$$

which is of the form:

$$Ax + By + Cz = 0$$

which is the equation of a plane surface.

When the equation has a positive result for particular xy values (pixel locations) then the xy location is on the forward facing side of the edge surface and when it has a negative value then the xy location is on the backward facing side of the surface. Thus, when all four equations representing the triangle in Figure 1 have a positive value then the pixel position is within the triangle as illustrated in Figure 2. This rule holds true for any shape used in preference to a triangle, e.g., a quadrilateral.

A preferred embodiment of the invention is shown in Figure 3. In this, there is a polygon setup unit 2 which receives vertex data defining triangles and supplies the facing surface data for the triangles to respective ones of the set of 32 surface processors 4. At the same time, for each triangle being processed by a surface processor 4 it supplies three sets of edge data to each one of three

SUBSTITUTE SHEET (RULE 26)

WO 00/28480

PCT/GB99/03704

- 5 -

arrays of edge processors 6. These each comprise a depth evaluation unit which determine whether or not the value for the edge surface they are processing is positive or negative for each of 32 particular pixel locations. The outputs of each of these is a positive or negative sign bit and these sign bits for the three surfaces is supplied to the appropriate surface processor for that triangle. When all of the sign bits are positive as described above, then that surface processor knows that the triangular surface it is processing is visible, that is to say it is not outside the edge of the triangle and thus, it will supply a depth value as an output which will go to a depth store, after which further tests can be made on it to determine whether or not it is to be used to make a contribution to the image being processed. If one of the sign bits is negative then the surface processor 4 does not need to do anything.

The edge processors operate in the x direction, i.e., along a scan line in an image and, in a system which uses an array of 32 surface processors 4, will typically operate in a tile based system processing blocks of 32 x 32 pixels. The input value to each edge processor will therefore be equivalent to $B_y + C$. The edge processor uses an inaccurate non-restoring division algorithm which operates on the edge of the triangle. This algorithm effectively calculates

$$x = \frac{C}{A}$$

This is possible because the y value is constant for a particular value of x and thus $B_y + C$ is a constant along a particular scan line.

Table 1 shows the arithmetic operation involved in calculating the position of a transition point from inside to outside (positive to negative depth) of an edge.

SUBSTITUTE SHEET (RULE 26)

WO 00/28480

PCT/GB99/03704

- 6 -

A	C	Operation
Shift A left 4 ($16 \times A$)	C	Add
Shift A left 3 ($8 \times A$)	$16A + C = C_1$	If $C_1 > 0$ and $A > 0$ then Subtract, otherwise Add
Shift A left 2 ($4 \times A$)	$C_1 + 8A = C_2$	If $C_2 > 0$ and $A = 0$ then Subtract, otherwise Add
Shift A left 1 ($2 \times A$)	$C_2 + 4A = C_3$	If $C_3 > 0$ and $A > 0$ then Subtract, otherwise Add
A	$C_3 + 2A = C_4$	If $C_4 > 0$ and $A > 0$ then Subtract, otherwise Add
A	$C_4 - A = C_5$	If $C_5 > 0$ and $A > 0$ then Subtract, otherwise C,
	$C_5 (-A) = C_6$	

Table 1 - Inaccurate Non-Restoring Division in an Edge Processor

The operation performed in stage 1A effectively moves the sample point to the middle in terms of x . This is possible because the setup unit moves the origin location $(x, y) = (0, 0)$ to the top lefthand corner of the tile. The operation column indicates the test performed to calculate whether an addition or subtraction should be performed on the accumulated C value in the next clock cycle. These tests are essentially a form of binary search, where each addition/subtraction moves us closer to the zero crossing point. For example, say that the 0 transition is at 13.

x location

Start	C = -ve	A = +ve	0
Add 16	C = +ve		16
Sub 8A	C = -ve		8
Add 4A	C = -ve		12
Add 2A	C = +ve		14
Sub A	C = 0 (+ve)		13
Sub A			12

The sign of the additions/subtractions which are performed by the edge processor are used to calculate the transition point or edge. Once this pixel position has been determined, it can be then used to create a mask for

SUBSTITUTE SHEET (RULE 26)

WO 00/28480

PCT/GB99/03704

- 7 -

a whole line of the tile. This mask represents a positive/negative depth value for each pixel within the line. The operation may be pipelined using the arrays of depth processors referred to above so that an edge mask for a line of pixels within a tile can be created every clock cycle. As explained above the y coefficient for the edge equation is accumulated into constant C before the edge is processed. This allows an edge mask for a complete tile of 32×32 pixels to be generated over 32 clock cycles where h is the height of the tile.

The masks for all three edges in the triangle are ANDed together to create a depth mask for the triangle. The signs of the accumulated depth at the pixel position is passed to the surface processors 4. When the depth is positive the surface is visible. Thus, using this method a triangle can be processed at the same speed as a single surface. Clearly, if four edge processors or more were available then quadrilaterals and other more complex shapes could be processed.

When the screen of the image is divided into a plurality of tiles the current hardware implementations require all objects within the scene to be processed for each tile. This is inefficient since it means that all the tiles have to be processed for all the objects.

In conventional rendering systems, rendering of the screen on a tile by tile basis requires objects to be clipped to tile boundaries, and therefore data defining the intersections with tile boundaries has to be defined (see Figure 4).

It is only necessary to process the objects which intersect with a particular region area. As explained above, if an object is defined in screen space, then a comparison of the vertices which define the object, such as a triangle, will yield a bounding box for that object. A bounding box defines a rectangular area within the screen which contains the object. Figure 4 shows a tiled

SUBSTITUTE SHEET (RULE 26)

WO 00/28480

PCT/GB99/03704

- 8 -

region of the screen with an object represented by a number of triangles within it.

A bounding box for a particular object can be aligned to tile boundaries so that a list of tiles within the bounding box can then be obtained. This list of tiles is a subset of all the tiles within the screen and approximates the tiles which intersect with the object. In the event that the bounding box with the object intersects with the whole of the screen area, then the object parameters (coordinates, shading data, etc.) are written into an area of memory within the system and a pointer to the start of the object data is generated.

The present rendering system operates on a tile by tile basis, processing the objects for each tile before progressing onto the subsequent one. The data structure is therefore used to identify the objects which must be processed for each tile. This is shown in Figure 5. In this, a list of tiles within the screen is created in a region or tile array 30. Each tile is defined by x and y limits. For each tile, a list of pointers to objects which must be processed for that tile is generated as an object list 32. There is a separate object list for each tile pointed to by the region array. The bounding box idea is used to create a list of tiles (with object lists) that the object pointer, which is created when data is written to memory, must be added to. However, the hardware needs to identify the tail of each object list so that an address for the object pointer to be written to can be derived. The most simple method of doing this is to store a tail pointer which points to the next free location on the list. This can be a header in the object list.

An enhancement of this is to use a cache which can be a smaller size. The cache stores a sub-set of the tail pointers. As an object will generally cross multiple tile boundaries, a miss upon the cache results in multiple tail pointers being read in and predicting the tiles which the

SUBSTITUTE SHEET (RULE 26)

WO 00/28480

PCT/GB99/03704

- 9 -

object traverses. This increases the efficiency of the cache. This also enables multiple images to be tiled at the same time by interleaving the object data and changing the cache contents. This switching involves storing the contents of the tail pointer cache, adjusting the area of memory linked to the cache and the area of memory used for storing objects. The effect of the context is now changed. That is to say, the cache is invalidated and a different set of data is now available to be tiled. To switch context back is the reverse operation and involves storage of the new context, the reversion of cache and object memory locations, and the invalidation of current cache.

The information for the object lists is now available. An address for the pointer which comes from the tail pointer cache and an object pointer pointing to an object which has a bounding box intersecting with that tile. All the object lists entries for the object being processed can then be written to memory and the next object processed.

This is implemented using the circuitry of Figure 10. In this, object data is received from the application program in the form of triangles, fans, strips and points. Initially the object data is all converted into strips, in a conversion unit 40. These are efficient in their memory usage. The converter 40 comprises a converter 42 for converting fans and faces to strips and converter 44 for points and lines to strips. The strip data is then provided to a bounding box generator 46 which calculates the bounding box for each triangle within the strip, and the bounding box for the whole strip. If the bounding box intersects with the screen area the object data is written to memory via local read/write arbiter 48, starting from the next available location. Otherwise the system moves on to the next strip. The address which this data is written to is passed down the pipeline.

SUBSTITUTE SHEET (RULE 26)

WO 00/28480

PCT/CB99/03704

- 10 -

A region generator 50 receives the bounding box information and generates a mask and tile identity for each tile within the bounding box for the whole strip. The tile identity is used to access a tail pointer cache 52 to read the next available pointer location. If this is the last address within the block, a new pointer block is allocated for this tile, and a link from the current block to the new one is generated.

A write request to a free address for the pointer, with the object address, and the mask for that object is placed into a queue. The tail pointer for the tile is then updated through the cache with the next available pointer. When there are sixteen entries in the write queue, the requests are sorted by page address, by the pointer sorter 54. These are written into the memory in a first access. This reduces the number of page breaks to the memory.

The most common type of cheap block RAM is DRAM. This is structured in pages and accesses which traverse pages. This is because there is a performance cost due to closing the current page and opening a new page. However, writing a pointer to the same object into multiple lists involves a large number of page transitions as each list may be on a different page. However, it is probable that there will be a similarity between one incoming object and the next object. This means that the next object is likely to be placed in similar object lists as current and previous objects. With an object list, the addresses are essentially sequential and it is therefore desirable to write as many pointers within the same list at the same time as there is address concurrency between pointers, this may be achieved by storing a number of pointers (e.g. over a range of objects) and sorting them into page groups before writing them to memory. This greatly reduces the number of page transitions and therefore increases the efficiency of the system.

SUBSTITUTE SHEET (RULE 26)

WO 00/28480

PCT/CB99/03704

- 11 -

For a given object data set for an image, it is not possible to determine the number of objects which will be visible in each tile. The worst case scenario is that it will be necessary to allocate enough memory for a pointer to every object for every tile. This would require a large amount of memory and increase the system cost. This can be reduced by allocating blocks for each object list and, when a block has been filled, allocating a new block and inserting a link to the new block. This means that the memory used is closer to the minimum amount required for object list storage. The size of a block will depend upon many factors such as the width of the memory and the available bandwidth.

In order to reduce both the number of object pointers and the size of object data further, another aspect of object coherency can be used. Since generally a group of triangles will be used to represent a larger object such as a tea-pot or sphere or animal, etc., there will be a large amount of commonality between triangles, i.e., the triangles will share vertices between them. By comparing the vertices against each other, it is possible to convert triangles to strips. A strip takes up less area of memory as only one or two vertices are required to define a new triangle and only one pointer is then required to point to all the objects in the strip. This reduces the number of object pointers even further and also reduces the amount of memory required, thereby resulting in an increase in efficiency in terms of memory and a performance increase due to bandwidth optimisations. In Figure 6 there is illustrated a triangle and a bounding box, this being the shaded portion. When it is processed using conventional methods, the region within which it falls covers a 5 x 5 array of tiles and it would be necessary to process it 25 times. However, if the image is first processed using a bounding box, to define the region which holds the range of x,y coordinates used by the triangle, it can be shown

SUBSTITUTE SHEET (RULE 26)

WO 00/28480

PCT/GB99/03704

- 12 -

that the triangle only needs to be processed 12 times,
i.e., it covers 12 tiles.

We have further appreciated that in fact the triangle
only falls within 10 of the tiles in the 4 x 3 array.

Thus reducing further the processing overhead.

Other examples of triangles which do not cover the
whole of the rectangular bounding box required to process
them are shown in Figures 7a-d. The most extreme example
of these is Figure 7d in which the triangle shown only in
fact falls in the 12 tiles illustrated in Figure 8. It
will be preferable to process only this set of tiles in
order to render that triangle.

The calculation of the minimal set of tiles to
represent a triangle begins with the crude rectangular
bounding box calculation. If the bounding box is only a
single tile in either height or width, there is clearly no
further optimisation that can be performed. Otherwise the
set will be reduced by consideration of each edge of the
triangle in turn.

Firstly, it is necessary to know whether the triangle
is defined by a clockwise or (cw) anti-clockwise (acw) set
of points. If this information is not available, it can
easily be calculated.

An edge can then be considered to be an infinitely
long line which divides the space into two halves. Sub-
spaces on either side of the edge are described as being
inside or outside the edge using the edge processors
described above with the inside sub-space being the one
that contains the triangle to which the edge belongs. The
triangle has its corners at the intersections of the edge
lines and the surface is the intersection of the inside
sub-spaces of the three edges.

Any tile that lies entirely on the outside of an edge
is not part of the minimal set because the triangle would
not be visible in that tile. If an edge is entirely
horizontal or vertical it need not be considered since all

SUBSTITUTE SHEET (RULE 26)

WO 00/28480

PCT/GB99/03704

- 13 -

the tiles in the rectangular bounding box already lie wholly or partly inside the edge.

In order to test whether a tile lies wholly on the outside of an edge, we need only test the point on that corner of the tile which is closest to the edge. If that point is on the outside of the edge, then we can be confident that the entire tile is also outside the edge. The position of this test point is determined by the orientation of the edge as indicated in the table given in Figure 9.

The edge itself can be described using the equation

$$y = mx + c$$

where x and y are coordinates of the screen, m represents the gradient of the line, and c is a constant. The valuation of $mx + c$ at the corner of a tile will give a value that is greater than, less than or equal to the y coordinate of that point. The comparison of the two values will indicate whether the point lies on the inside or outside of the edge. The interpretation of this result depends on the orientation of the edge as given in the table in Figure 9.

For each edge of the triangle, each tile in the rectangular bounding box must be processed in this way to decide whether or not it should be excluded from the minimal set.

It should be noted that the test point at the corner of a tile is also the test point for a larger rectangular set of tiles. In Figure 10, knowing that the tile marked where the test point is outside an edge means that all the shaded tiles must also be outside that edge. In this example, where the test point is at the bottom right, it is most efficient to process the tiles of the rectangular bounding box from right to left and from bottom to top, in all there is a large number of tiles may be excluded from the minimal set with the minimum number of tests. When

SUBSTITUTE SHEET (RULE 26)

WO 00/28480

PCT/GB99/03704

- 14 -

the test point is in a different corner of the tile, the order of processing would be changed accordingly.

SUBSTITUTE SHEET (RULE 26)

WO 00/28480

PCT/G899/03704

- 15 -

CLAIMS

1. A method for shading 3-dimensional computer generated images comprising the steps of:

representing each object in the image by a set
5 of flat polygons;

for each polygon, supplying a set of vertex data defining the vertices of the polygon, along with data defining the orientation of its surface;

10 computing edge data from the vertex data, the edge data being equivalent to data defining a surface perpendicular to a viewpoint and facing towards the polygon;

for each pixel used to view the portion of the image in which the polygon is located, deriving a depth
15 value for each surface; and

shading that pixel when the depth values indicate that the polygon is visible at that pixel.

2. A method according to claim 1 in which a surface is visible at a particular pixel when its depth
20 value is positive value and each of the edge surfaces have positive depth values at that pixel.

3. A method according to claim 1 including the step of determining edge transitions from the edge data, and deriving an edge mask for each edge.

25 4. A method according to claim 3 in which the edge mask is derived line by line of pixels, one line per clock cycle.

5. A method according to claim 4 in which the edge mask is derived line by line for each of a plurality of
30 rectangular sub-regions of the image.

SUBSTITUTE SHEET (RULE 26)

WO 00/28480

PCT/GB99/03704

- 16 -

6. A method according to claim 4 or 5 in which the edge mask is derived for each surface and includes the step of deriving a depth mask for that surface from the edge masks.

7. Apparatus for shading 3-dimensional computer generated images comprising:-
means for representing each object in the image by a set of flat polygons;

means for supplying a set of vertex data defining the vertices of each polygon, along with data defining the orientation of its surface;

means for computing edge data from the vertex data, the edge data being equivalent to data defining a surface perpendicular to a viewpoint and facing towards the polygon;

means for deriving depth data for each surface for each pixel used to view the portion of the image in which the polygon is located; and

means for shading that pixel when the depth values indicate that the polygon is visible at that pixel.

8. Apparatus according to claim 7 in which a surface is visible at a particular pixel when its depth value is positive and each of the edge surfaces surrounding it has a positive depth value at that pixel.

9. Apparatus according to claim 7 including means for determining edge transitions from the edge data, and means for deriving an edge mask for each edge from the edge transitions.

10. Apparatus according to claim 9 in which the edge mask is derived line by line of pixels in the image, one line per clock cycle.

SUBSTITUTE SHEET (RULE 26)

WO 00/28480

PCT/GB99/03704

- 17 -

11. Apparatus according to claim 10 in which the edge mask is derived line by line for each of a plurality of rectangular sub-regions of the image.

5 12. Apparatus according to claim 10 or 11 in which the edge mask is derived for each edge defining a surface and further comprising means for deriving a depth mask for that surface from the edge mask.

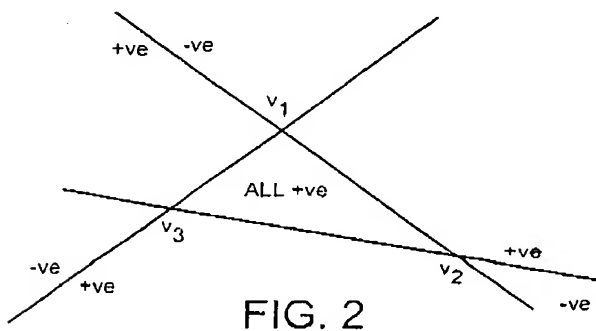
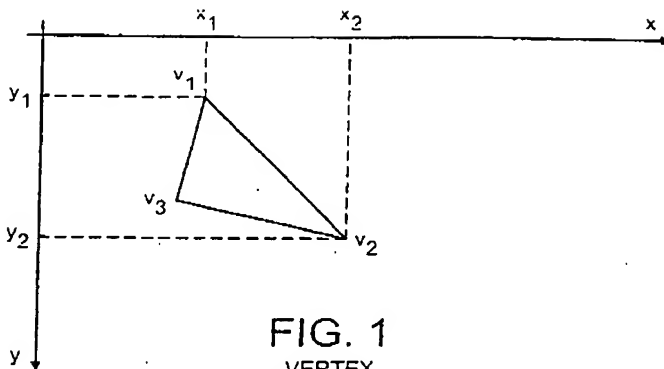
10 13. A method according to any of claims 1 to 6 in which no edge clipping is required when a polygon intersects a boundary of the image being shaded.

SUBSTITUTE SHEET (RULE 26)

WO 00/28480

PCT/GB99/03704

1 / 6



SUBSTITUTE SHEET (RULE 26)

WO 00/28480

PCT/GB99/03704

2 / 6

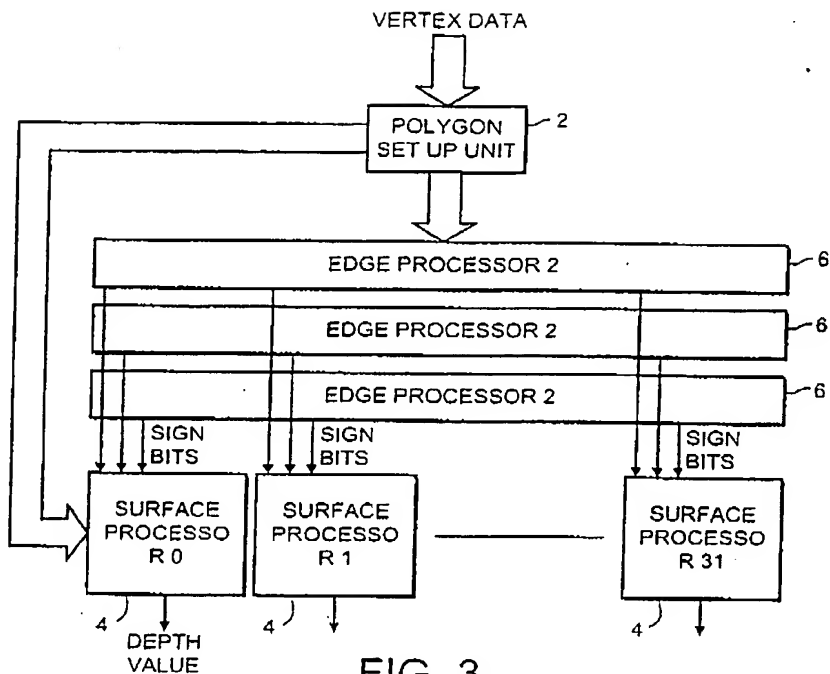


FIG. 3

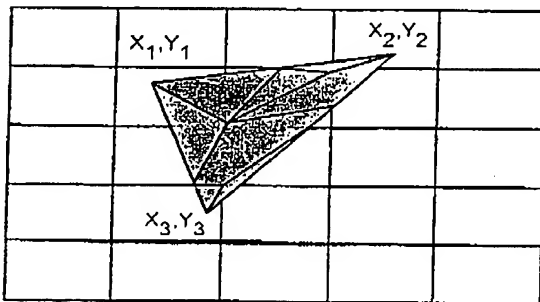


FIG. 4

SUBSTITUTE SHEET (RULE 26)

WO 00/28480

PCT/GB99/03704

3 / 6

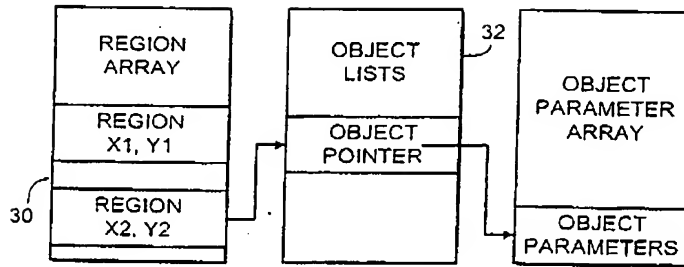


FIG. 5

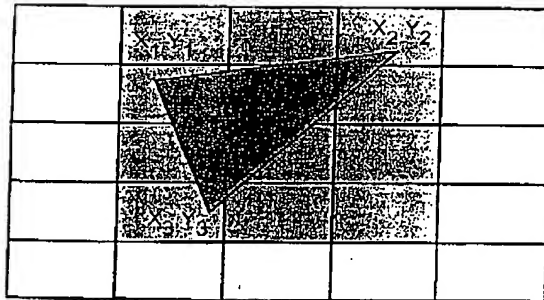


FIG. 6

SUBSTITUTE SHEET (RULE 26)

WO 00/28480

PCT/GB99/03704

4 / 6

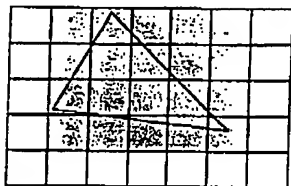


FIG. 7a

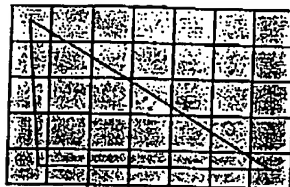


FIG. 7b

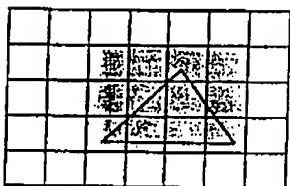


FIG. 7c

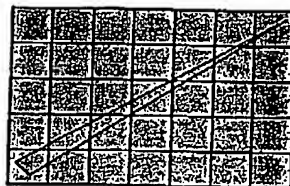


FIG. 7d

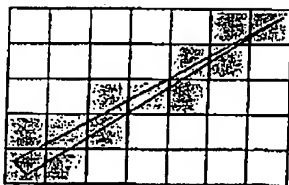


FIG. 8



FIG. 10

SUBSTITUTE SHEET (RULE 26)

WO 00/28480

PCT/GB99/03704

5 / 6

TABLE 1












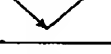

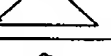


TEST POINT	TRIANGLE	EDGE ORIENTATION	CONDITION FOR TEST POINT ON OUTSIDE OF EDGE.
	CW		$y < mx + c$
	CW		$y < mx + c$
	CW		$y > mx + c$
	CW		$y > mx + c$
	ACW		$y > mx + c$
	ACW		$y > mx + c$
	ACW		$y < mx + c$
	ACW		$y < mx + c$

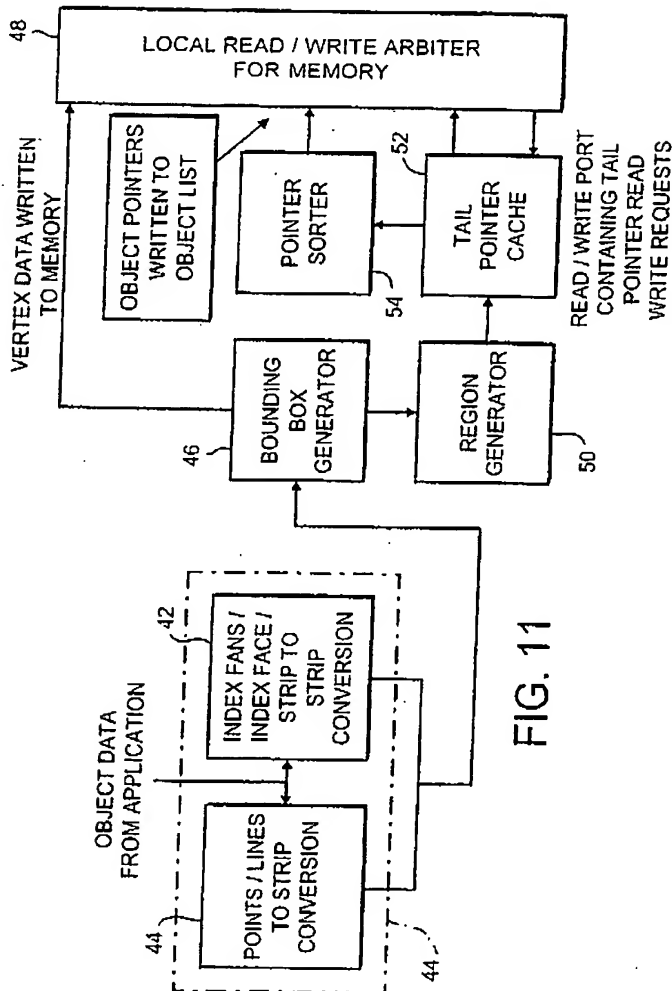
FIG. 9

SUBSTITUTE SHEET (RULE 26)

WO 00/28480

PCT/GB99/03704

6 / 6



SUBSTITUTE SHEET (RULE 26)

INTERNATIONAL SEARCH REPORT

Intern. Application No.
PCT/GB 99/03704A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 G06T15/40

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 G06T

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	EP 0 758 113 A (MITSUBISHI ELECTRIC CORP) 12 February 1997 (1997-02-12) column 6, line 41 - line 52 column 10, line 13 - line 58 ----	1-3,7-9
Y	EP 0 549 183 A (GEN ELECTRIC) 30 June 1993 (1993-06-30) column 4, line 42 - column 5, line 32 column 6, line 19 - line 34 -----	1-3,7-9

☐ Further documents are cited in the continuation of box C☒ Patent family members are listed in annex

Special categories of cited documents

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another claim or other special reason (as specified)
- *J* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
- *G* document member of the same patent family

Date of the actual completion of the international search

10 February 2000

Date of mailing of the international search report

17/02/2000

Name and mailing address of the ISA

European Patent Office, P. O. Box 1 Patentplaza 2
NL - 2280 HV Rijswijk
Tel: (+31-70) 340-2040 Te: 31 651 epo nl
Fax: (+31-70) 340-2016

Authorized officer

Burgaud, C

Form PCT/ISA/210 (40000 sheets) (July 1992)

INTERNATIONAL SEARCH REPORT

information on patent family members

Intern. Application No
PCT/GB 99/03704

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 0758118 A	12-02-1997	JP 9050537 A	18-02-1997
		US 5831623 A	03-11-1998
EP 0549183 A	30-06-1993	JP 2809955 B	15-10-1998
		JP 5274446 A	22-10-1993

Form PCT/ISA/210 (patent family annex) (July 2002)

Dec-02-2004 02:59pm From:8586456880

+

T-673 P.031/053 F-509